GAME THEORY: ANALYSIS OF STRATEGIC THINKING.
Bayes Rule and Signaling

Pierpaolo Battigalli Bocconi University

A.Y. 2006-07

1 Bayes Rule

Consider an individual i who is initially uncertain about two things: a parameter θ, and a variable (or a vector of variables) x. In a second stage i observes x.

Assume $x \in X$ and $\theta \in \Theta$, with X and Θ finite. Even though i cannot observe θ, he is able to assess the conditional probability of each value $x \in X$ conditional on each $\theta \in \Theta$, $\Pr[x|\theta]$.

For example, consider an urn of unknown composition. It is only known that the urn contains between 1 and 10 balls, which may differ only in the color, Black or White. A first ball is drawn, its color is observed and then it is put back into the urn. Then a second ball is drawn (possibly the same as before) and its color observed. Let θ denote the proportion of White balls in the urn. The set Θ of possible values of this parameter is finite:

$$\Theta = \bigcup_{n=1}^{10} \left\{ \theta : \theta = \frac{k}{n} \text{ for some } k = 0, ..., n \right\}$$

Let Black correspond to number 0 (since black is the absence of light) and White correspond to number 1. Then X can be identified with the set of all ordered pairs (x_1, x_2) with $x_k \in \{0,1\}$, that is, $X = \{0,1\}^2$. The probability of $x = (x_1, x_2)$ conditional on the proportion of white balls being θ is $\Pr[x|\theta] = \theta^{x_1}(1-\theta)^{x_1}\theta^{x_2}(1-\theta)^{1-x_2}$. For example, $\Pr[(1,0)|\theta] = \theta(1-\theta)$.

Mister i also assigns a probability $\Pr[\theta]$ to all the possible values of θ. Note that $\Pr[x|\theta]$ is well defined even if i assigns probability 0 to θ. For example, for some reason i may be certain that the urn does not contain more than 5 balls, and hence $\Pr[\theta = \frac{1}{10}] = 0$. Yet i thinks that if θ *were* $\frac{1}{10}$ then the probability of $x = (0,0)$ *would be* $\frac{9}{10} \times \frac{9}{10}$, that is, $\Pr[(0,0)|\frac{1}{10}] = \frac{81}{100}$.

The law of conditioning says that for any two events E and F,

$$\text{if } \Pr[F] > 0, \text{ then} \Pr[E|F] = \frac{\Pr[E \cap F]}{\Pr[F]}.$$

Note that the same condition can be written more compactly as

$$\Pr[E \cap F] = \Pr[E|F] \Pr[F].$$

Here we may assume that the space of uncertainty is $\Omega = \Theta \times X$. Each x corresponds to event $E = \Theta \times \{x\}$, each θ corresponds to event $F = \{\theta\} \times X$, each pair (θ, x) is the singleton $\{(\theta, x)\}$. Then the joint probability of the pair (θ, x) is

$$\Pr[(\theta, x)] = \Pr[x|\theta] \Pr[\theta].$$

Note that the collection of events $\{F \subseteq \Omega : F = \{\theta\} \times X \text{ for some } \theta \in \Theta\}$ forms a partition of $\Omega = \Theta \times X$. This partition of Ω induces a corresponding partition of every event $E \subseteq \Omega$, and in particular of events of the form $E = \Theta \times \{x\}$. Then we can compute probability of any x (event $\Theta \times \{x\}$) starting from the probabilities of the "cells" $\{(\theta, x)\} = (\Theta \times \{x\}) \cap (\{\theta\} \times X)$ with $\theta \in \Theta$. In turn, these probabilities can be obtained from the conditional and prior probabilities $\Pr[x|\theta]$ and $\Pr[\theta]$, $\theta \in \Theta$. Thus, we obtain the formula expressing the marginal, or **total probability** of x as

$$\Pr[x] = \sum_{\theta' \in \Theta} \Pr[(x, \theta')] = \sum_{\theta' \in \Theta} \Pr[x|\theta'] \Pr[\theta']. \tag{TotP}$$

The problem is to derive from these elements $\Pr[\theta|x]$, the probability that i would assign to each $\theta \in \Theta$ upon observing any $x \in X$. There are two possibilities, either $\Pr[x] = 0$ or $\Pr[x] > 0$.

If $\Pr[x] = 0$, $\Pr[\theta|x]$ cannot be derived from the previous data. *This does not mean that i is unable to assess the conditional probability* $\Pr[\theta|x]$, it only means that $\Pr[\theta|x]$ is not determined by the other probabilistic assessments expressed above.

If $\Pr[x] > 0$, then $\Pr[\theta|x] = \frac{\Pr[(\theta,x)]}{\Pr[x]}$. Substituting $\Pr[x]$ with the expression given by (TotP) we obtain

$$\Pr[\theta|x] = \frac{\Pr[(\theta, x)]}{\sum_{\theta' \in \Theta} \Pr[x|\theta'] \Pr[\theta']}, \tag{BFor}$$

(BFor) is known as **Bayes Formula,** that is, an equation expressing $\Pr[\theta|x]$ as a function of the conditional probabilities $\Pr[x|\theta']$ and the prior probabilities $\Pr[\theta']$, with $\theta' \in \Theta$.

Bayes Rule says that whenever (BFor) *can* be applied, then it *must* be applied. Since (BFor) can be applied if and only if $\Pr[x] > 0$, we may write Bayes Rule in the following compact form:

$$\forall x \in X, \forall \theta \in \Theta,\ \Pr[\theta|x] \left(\sum_{\theta' \in \Theta} \Pr[x|\theta'] \Pr[\theta'] \right) = \Pr[(\theta, x)]. \tag{BRule}$$

Bayes Rule is *not violated* if either $\Pr[x] = 0$ (both sides of (BRule) are zero), or $\Pr[x] > 0$ and $\Pr[\theta|x]$ is computed with (BFor). Therefore, whenever an assessment of prior and conditional probabilities satisfies (BRule), we say that it is *consistent with Bayes rule*. Note that Bayes rule "holds" even if $\Pr[x] = 0$ (another way to express this point is that if the antecedent in the material implication $\Pr[x] > 0 \Rightarrow \frac{\Pr[(\theta,x)]}{\sum_{\theta' \in \Theta} \Pr[x|\theta'] \Pr[\theta']}$ is false, the material implication holds).[1]

2 Signaling Games and Perfect Bayesian Equilibrium

Now assume that θ is known to player 1 (female), whose action[2] $a_1 \in A_1$ is observed by player 2 (male) before he chooses $a_2 \in A_2(a_1)$. The possible parameter values $\theta \in \Theta$ are called **types** of player 1, the informed player. The actions of of the informed player are also called **signals** or **messages** because they may reveal her private information. Let $A_2 = \bigcup_{a_1 \in A_1} A_2(a_1)$. The payoffs of players 1 and 2 are given by the functions

$$\begin{aligned} u_1 &: \Theta \times A_1 \times A_2 \to \mathbb{R}, \\ u_2 &: \Theta \times A_1 \times A_2 \to \mathbb{R}. \end{aligned}$$

A **behavior strategy** for player 1 is an array of probability measures $\beta_1 = (\beta_1(\cdot|\theta))_{\theta \in \Theta} \in \prod_{\theta \in \Theta} \Delta(A_1) = [\Delta(A_1)]^{\Theta}$. A behavior strategy for player 2 is an array of probability measures $\beta_2 = \beta_2(\cdot|a_1))_{a_1 \in A_1} \in \prod_{a_1 \in A_1} \Delta(A_2(a_1))$.

Here player 2 has the role of individual i in the previous section, with $X = A_1$. Player 2 is initially uncertain about (θ, a_1) and has a **prior** probability measure $\rho \in \Delta(\Theta)$. We assume for simplicity that $\forall \theta \in \Theta,\ \rho(\theta) > 0$. This *prior* is an *exogenously given* element of the model, whereas β_1 and β_2 are *endogenous*, i.e. they have to be determined through equilibrium analysis.

In equilibrium, β_1 represents the assessment of player 2 about the probability of each action of player 1 conditional on each possible parameter value θ. Thus, the probabilistic assessment of player 2 is such that

$$\begin{aligned} \forall \theta &\in \Theta,\ \Pr[\theta] = \rho(\theta), \\ \forall \theta &\in \Theta, \forall a_1 \in A_1,\ \Pr[a_1|\theta] = \beta_1(a_1|\theta),\ \Pr[(\theta, a_1)] = \beta_1(a_1|\theta)\rho(\theta), \\ \forall a_1 &\in A_1,\ \Pr[a_1] = \sum_{\theta' \in \Theta} \beta_1(a_1|\theta')\rho(\theta'). \end{aligned}$$

We let $\mu(\theta|a_1)$ denote the probability that player 2 would assign to θ upon observing a_1. Since player 2 chooses a_2 after he has observed a_1 in order to maximize the expectation of $u_2(\theta, a_1, a_2)$, the system of conditional probabilities

[1] The material implication $p \Rightarrow q$ is verified if either p is false, or both p and q are true.

[2] In some models the set of feasible actions of player 1 depends on θ and is denoted by $A_1(\theta)$. The set of potentially feasible actions of player 1 is $A_1 = \bigcup_{\theta \in \Theta} A_1(\theta)$.

$\mu = (\mu(\cdot|a_1))_{a_1 \in A_1}$ is an essential ingredient of equilibrium analysis. In the technical language of game theory μ is called **"system of beliefs"**.

Now, if $\Pr[a_1] = \sum_{\theta' \in \Theta} \beta_1(a_1|\theta')\rho(\theta') > 0$ then Bayes formula applies and

$$\mu(\theta|a_1) = \frac{\Pr[(\theta, a_1)]}{\Pr[a_1]} = \frac{\beta_1(a_1|\theta)\rho(\theta)}{\sum_{\theta' \in \Theta} \beta_1(a_1|\theta')\rho(\theta')}.$$

Since β_1 is endogenous, also the system of beliefs μ *is endogenous*. Thus we have to determine through equilibrium analysis the triple (β_1, β_2, μ). In the technical game-theoretic language (β_1, β_2, μ) (a profile of behavioral strategies plus a system of beliefs) is called **"assessment".**

For any given assessment (β_1, β_2, μ) we use the following notation to abbreviate conditional expected payoff formulas:

$$\begin{aligned}
\mathrm{E}^{\beta_2, a_1}[u_1|\theta] \quad &: \quad = \sum_{a_2 \in A_2(a_1)} \beta_2(a_2|a_1) u_1(\theta, a_1, a_2), \\
\mathrm{E}^{\mu, a_2}[u_2|a_1] \quad &: \quad = \sum_{\theta \in \Theta} \mu(\theta|a_1) u_2(\theta, a_1, a_2).
\end{aligned}$$

Thus, $\mathrm{E}^{\beta_2, a_1}[u_1|\theta]$ is the expected payoff for player 1 of choosing action a_1 given that her type is θ and assuming that her "conjecture" about the behavior of player 2 is represented by β_2.[3] Similarly, $\mathrm{E}^{\mu, a_2}[u_2|a_1]$ is the expected payoff for player 2 of choosing action a_2 given that he has observed a_1 and assuming the her conditional beliefs about θ are represented by μ.

In equilibrium, an action of player i can have positive (conditional) probability only if it maximizes the (conditional) expected payoff of i. Furthermore, the equilibrium assessment must be consistent with Bayes rule. Therefore we obtain three equilibrium conditions for the three "unknowns" (β_1, β_2, μ):

Definition. Assessment (β_1, β_2, μ) is a *perfect Bayesian equilibrium (PBE)* if satisfies the following conditions:

$$\forall \theta \in \Theta, \mathrm{Supp}\beta_1(\cdot|\theta) \subseteq \arg\max_{a_1 \in A_1} \mathrm{E}^{\beta_2, a_1}[u_1|\theta] \qquad (\mathrm{BR}_1)$$

$$\forall a_1 \in A_1, \mathrm{Supp}\beta_2(\cdot|a_1) \subseteq \arg\max_{a_2 \in A_2(a_1)} \mathrm{E}^{\mu, a_2}[u_2|a_1] \qquad (\mathrm{BR}_2)$$

$$\forall a_1 \in A_1, \forall \theta \in \Theta, \mu(\theta|a_1) \sum_{\theta' \in \Theta} \beta_1(a_1|\theta')\rho(\theta') = \beta_1(a_1|\theta)\rho(\theta). \qquad (\mathrm{CONS})$$

Clearly (CONS), consistency with Bayes rule, can also be expressed as

$$\begin{aligned}
\forall a_1 \quad &\in \quad A_1, \forall \theta \in \Theta, \\
\sum_{\theta' \in \Theta} \beta_1(a_1|\theta')\rho(\theta') \quad &> \quad 0 \Rightarrow \mu(\theta|a_1) = \frac{\beta_1(a_1|\theta)\rho(\theta)}{\sum_{\theta' \in \Theta} \beta_1(a_1|\theta')\rho(\theta')}.
\end{aligned}$$

[3] There is no loss of generality in representing a conjecture of player 1 as a behavioral strategy of player 2. If player 1 had a conjecture of the form $\sigma_2 \in \Delta(S_2)$, where S_2 is the set of pure strategies of player 2, then we could derive from σ_2 a realization-equivalent behavioral strategy. A similar argument holds for the conjecture of player 2 about player 1.

Note that each equilibrium condition involves two out of the three vectors of endogenous variables β_1, β_2 and μ: (BR_1) says that that each mixed action $\beta_1(\cdot|\theta) \in \Delta(A_1)$ ($\theta \in \Theta$) is a best reply to β_2, (BR_2) says that each mixed action $\beta_2(\cdot|a_1) \in \Delta(A_2(a_1))$ ($a_1 \in A_1$) is a best reply to the conditional belief $\mu(\cdot|a_1) \in \Delta(\Theta)$, and (CONS) says that β_1 and μ (together with the exogenous prior ρ) are consistent with Bayes rule.

It should be emphasized that, for some action a_1, $\Pr[a_1] = \sum_{\theta' \in \Theta} \beta_1(a_1|\theta')\rho(\theta')$ *may be zero.* For example, suppose that for each $\theta \in \Theta$, there is some action $a_1(\theta)$ such that $\mathrm{E}^{\beta_2, a_1(\theta)}[u_1|\theta] > \mathrm{E}^{\beta_2, a_1}[u_1|\theta]$. Then action/message a_1 must have zero probability because it is not a best reply for any θ. *Yet we assume that the belief* $\mu(\cdot|a_1)$ *is well defined* and player 2 takes a best reply to this belief. This is a *perfection* requirement analogous to the subgame perfection condition for games with observable actions and complete information. A perfect Bayesian equilibrium satisfies perfection and consistency with Bayes rule.

Furthermore, even if $\mu(\cdot|a_1)$ cannot be computed with Bayes formula, it may still be the case that the equilibrium conditions put constraints on the possible values of $\mu(\cdot|a_1)$. The following example illustrates this point.

$$
\begin{array}{ccccccccccc}
 & & & & & & & & & & (\mathbf{0},3)\\
 & & l & & \{\frac{1}{2}\} & & r & & & \nearrow & \\
(\mathbf{1},1) & \longleftarrow & - & - & 1 & - & - & \longrightarrow & 2 & & \\
 & & & & \theta' & & & & \vdots & \searrow & \\
 & & & & & & & & \vdots & & (\mathbf{0},0)\\
 & & & & & & & & \vdots & & \\
 & & & & & & & & \vdots & & (\mathbf{2},0)\\
 & & & & \theta'' & & & & \vdots & \nearrow & \\
(\mathbf{1},1) & \longleftarrow & - & - & 1 & - & - & \longrightarrow & 2 & & \\
 & & l & & \{\frac{1}{2}\} & & r & & & \searrow & \\
 & & & & & & & & & & (\mathbf{0},1)
\end{array}
$$

The payoffs of player 1 are in **bold.** $A_2(l)$ is a singleton and therefore the action of player 2 after l is not shown. If player 1 goes right (r) then player 2 can go up (u) or down (d), i.e. $A_2(r) = \{u, d\}$.

Note that action r is dominated for type θ'. Therefore $\beta_1(r|\theta') = 0$ in every PBE. Now we show that in equilibrium we also have $\beta_1(r|\theta'') = 0$. Suppose, by way of contradiction, that $\beta_1(r|\theta'') > 0$. Then Bayes formula applies and $\mu(\theta''|r) = 1$. But then the best reply of player 2 is down, $\beta_2(d|r) = 1$, and the best reply of type θ'' is left, $\beta_1(l|\theta'') = 1 - \beta_1(r|\theta'') = 1$, contradicting our initial assumption.

We conclude that in every PBE r is chosen with probability zero and $\mu(\cdot|r)$ cannot be determined with Bayes formula. Yet the equilibrium conditions put a constraint on $\mu(\cdot|r)$: in equilibrium d must be (weakly) preferred to u (if player 2 chooses u after r then type θ'' chooses r and we have just shown that this cannot happen in equilibrium). Therefore

$$
\begin{aligned}
\mu(\theta''|r) &\geq 3\mu(\theta'|r) \\
\text{or } \mu(\theta''|r) &\geq \frac{3}{4}.
\end{aligned}
$$

The set of equilibrium assessments is

$$
\begin{aligned}
&\left\{(\beta_1, \beta_2, \mu) : \beta_1(l|\theta') = \beta_1(l|\theta'') = 1, \beta_2(d|r) = 1, \mu(\theta''|r) > \frac{3}{4}\right\} \\
&\cup \left\{(\beta_1, \beta_2, \mu) : \beta_1(l|\theta') = \beta_1(l|\theta'') = 1, \beta_2(d|r) \geq \frac{1}{2}, \mu(\theta''|r) = \frac{3}{4}\right\}.
\end{aligned}
$$

These assessments are examples of "pooling" equilibria. A **pooling equilibrium** is a PBE assessment where all types of player 1 choose the same pure action with probability one: there exists $a_1^* \in A_1$ such that $\forall \theta \in \Theta$, $\beta_1(a_1^*|\theta) = 1$. In this case Bayes rule implies that the posterior on θ conditional on the equilibrium action a_1^* is the same as the prior: $\mu(\cdot|a_1^*) = \rho(\cdot)$.

The polar case is when different types choose different pure actions: a **separating equilibrium** is a PBE assessment such that each type θ of player 1 chooses some action $a_1(\theta)$ with probability one ($\beta_1(a_1(\theta)|\theta) = 1$) and $a_1(\theta') \neq a_1(\theta'')$ for all θ' and θ'' with $\theta' \neq \theta''$. A separating equilibrium may exist only if A_1 has at least as many elements as Θ. *If A_1 and Θ have the same number of elements* (cardinality) then in a separating equilibrium each action is chosen with positive probability (because $\rho(\theta) > 0$ for each $\theta \in \Theta$) and the action of player 1 *perfectly reveals* her private information (if A_1 has more elements than Θ then the actions that are chosen by some type are perfectly revealing, the others need not be revealing).

The following signaling game provides an example of separating equilibrium (the payoffs of the informed player are in **bold,** call the downward action of player 2 a and the upward action f):

$$
\begin{array}{ccccccccccc}
(\mathbf{1},0) &&&&&&&&&& (\mathbf{0},0) \\
& \nwarrow &&& s & \{\frac{9}{10}\} & w &&& \nearrow & \\
&& 2 & \longleftarrow & - \quad - & 1 & - \quad - & \longrightarrow & 2 && \\
& \swarrow & \vdots &&& \theta^s &&& \vdots & \searrow & \\
(\mathbf{2},1) && \vdots &&&&&& \vdots && (\mathbf{1},1) \\
&& \vdots &&&&&& \vdots && \\
&& \vdots &&&&&& \vdots && \\
(\mathbf{0},1) && \vdots &&&&&& \vdots && (\mathbf{1},1) \\
& \nwarrow & \vdots &&& \theta^w &&& \vdots & \nearrow & \\
&& 2 & \longleftarrow & - \quad - & 1 & - \quad - & \longrightarrow & 2 && \\
& \swarrow &&& s & \{\frac{1}{10}\} & w &&& \searrow & \\
(\mathbf{1},0) &&&&&&&&&& (\mathbf{2},0)
\end{array}
$$

The game can be interpreted as follows: a truck driver (player 1) enters in a pub where an aggressive customer (player 2) has to decide whether to start a a fight (f, upward) or acquiesce (a, downward). There are two types of truck drivers: 90% of them are surly (θ^s) and like to eat sausages (s) for breakfast; the remaining 10% are wimps (θ^w) and prefer a dessert with whipped cream (w). Each type of truck driver receives a utility equal to 2 from her favorite breakfast and a utility equal to 1 from the other breakfast. Furthermore both types incur a loss of 1 util if they have to fight. Player 2 prefers to fight with a wimp and to avoid the fight with a surly driver.

This game has only one "reasonable" PBE and it is separating.[4] Indeed, the payoffs are such that for each type of driver it is weakly dominant to have her preferred breakfast thus iterated deletion of weakly dominated actions yields the equilibrium $\beta_1(s|\theta^s) = 1 = \beta_1(w|\theta^w)$, $\beta_2(a|s) = 1 = \beta_2(f|w)$, $\mu(\theta^s|s) = 1 = \mu(\theta^w|w)$.

The game has also two sets of pooling equilibria (meaning that one type of player 1 chooses a weakly dominated action). In the first set of assessments each type has sausages for breakfast and player 2 would fight if and only he observed a whipped-cream breakfast: $\beta_1(s|\theta^s) = 1 = \beta_1(s|\theta^w)$, $\beta_2(a|s) = 1 = \beta_2(f|w)$, $\mu(\theta^s|s) = \frac{9}{10}$, $\mu(\theta^w|w) \geq \frac{1}{2}$. In the second set of assessments each type has whipped cream for breakfast and player 2 would fight if and only if he observed a sausage breakfast: $\beta_1(w|\theta^s) = 1 = \beta_1(w|\theta^w)$, $\beta_2(a|w) = 1 = \beta_2(f|s)$, $\mu(\theta^s|w) = \frac{9}{10}$, $\mu(\theta^w|s) \geq \frac{1}{2}$.

[4] Actually, this example is a modification of a well-known game where the cost of a fight is *larger* than the marginal benefit from having the preferred breakfast and all equilibria are pooling.